

FMS Post 041917.txt

I just got back from a hearing in a court case involving Freenet. With the exception of a few (3?) documents that are still sealed, the government's documentation on the attack is now public record, as is the hearing itself. The transcript has not yet been published, (it is expected to take two weeks) but I took notes.

The earlier Reddit discussion thread about the paper on a Missouri police department attack that leaked due to a misconfigured Sharepoint instance is here:

https://www.reddit.com/r/Freenet/comments/4ebw9w/more_information_on_law_enforcements_freenet/

With that in mind, as a matter of public record, I can confirm:

- * In September 2011, Special Investigator Wayne Becker with the Missouri ICAC Task Force started collecting publicly available keys of child pornography. Initially the investigation only considered top-level manifest blocks, but when they found that to be ineffective they fetched the splitfile and considered all blocks of the file. That has been expanded into an automated process which runs daily and scrapes keys from Frost child porn boards. As of the last time SI Becker checked, there are 75k manifest blocks and 170 million split keys in this database. They call these Files of Interest.

- * Starting April 2012, law enforcement has been running modified Freenet nodes which connect to opennet. Initially this was anywhere from 1 to 8 nodes, they logged `_all` requests and inserts_ they observed to CSVs, and ran through most of 2014. Those logs have been retained. This was with a patch developed by people at the University of Massachusetts Amherst.

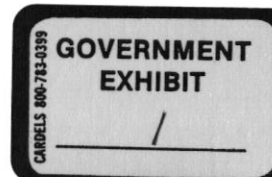
- * In winter 2014, the university researchers developed an improved version which only logs requests for blocks from Files of Interest. It can optionally, defaulting to disabled, forge DataNotFound results for FOI blocks. There are around 30 of these nodes running, and have been for two years now.

- * When the attack produces results which indicate, in the police's judgement, that a node is probably requesting child pornography, they will fetch the file in question, verify its contents, and retain the file.

- * Dr. Brian Levine is part of the team that developed the improved attack. He made an analogy about distributing M&Ms among a room of people that went something like:

- > One person starts with a bag of 100 M&Ms, someone chooses 4 people to give them to, split roughly evenly between. Each of them do the same until each person has M&Ms. Someone receiving 25, if they know the starting M&M count and the number of people the person giving them M&Ms is splitting their M&Ms between, can evaluate the probability that the person giving M&Ms is the one holding the bag.

Here an M&M is a request for a block, and the attack allows making a prediction as to whether the node that sent the request is requesting the entire file or just relaying requests. To do this it uses knowledge of the total size of the file in question from the collection of keys, and the peer count of the peer - shared



FMS Post 041917.txt

by default. ("Shall we send our peers' locations to our peers? Doing so helps routing but gives some information away to a potential attacker." under Config > Core in advanced mode.) It compares the received request count to a model which assumes uniform request distribution. They ran simulations, and claim to have established a false positive rate of 2% because they ran it against HTL 16 requests and it gave a positive result 2% of the time. (Even though an HTL 16 request indicates the node sending the request is never the originator given default settings.) I am not convinced this is a valid way to establish a false positive rate. These requests are one hop more diffuse, so they are not representative of the level of diffusion present in the actual requests they use it on. I'm not clear on the distribution of distances requests can be expected to have been probably routed for a given HTL; that would be relevant here.

* Using this technique they have performed, either with a search warrant or with consent, over 50 searches in the US and Canada.

Here's a document I used as a source for some of this: an Affidavit Referral Narrative. [0] (Search count from page 2 #13.) I redacted the section on specifics to the case out of respect for the defendant, but the broad strokes are:

1. An "undercover _Freenet_ node" routes requests from a node for blocks from an FOI, and logs them.
2. Later analysis on logged requests suggests the peer in question was the originator.
3. Through an IP geolocation check (for jurisdiction) and an ISP subpoena the police obtain a name and address.
4. Police file for a warrant, telling the judge the "number and timing of the requests was significant enough to indicate that the IP address was the apparent original requester of the file." It is my opinion that the warrant affidavit does not give enough information for a judge to independently assess the reasonability of this conclusion.
5. Knock knock.

If people care I can go through and redact the things specific to the defendant more precisely; it was easier to draw two big rectangles. (I might also be persuaded to bother to figure out how to redact PDFs properly, but for the time being I'm trying to get this posted quickly and imagemagick is easy.)

The defense will move to unseal one of the documents. For those who have PACER access, the case number is `4:16-cr-258 CEJ (NAB)`. Also it turns out some cars' black boxes record GPS location information.

[0]

CHK@06eT-w4owv1SdjfiwBrb5Ses0IPcWcdElgPE6aKxwAM,7wL-tcju7FkwbkGuFTC16uYy1cWxcLH1HTonhoIvYws,AAMC--8/narrativ

FMS Post 041917.txt

e-0.png

CHK@R77s92~0zfjMD6yAVd92~y9kka5lBHaf8LVpFK5HEEg,0352b5ygyrd7K-CAVtrVzCRx3hKZ6U61ogp35oTPB2A,AAMC--8/narrativ

e-1.png

CHK@jMPyWLeUQztSHT5cPZW48goy9xwdIwTmXasc8QkdDZ8,zHajD4SkJlZubYogzXs8hRWcuDSEVhRC38B0nkH47dQ,AAMC--8/narrativ

e-2.png

CHK@RIK1x8aYc2psKkvx~Se4EAbvttYjQTjq2g2M10SmMyc,rhT8AD~gsiW0kWEI91rcHxVeBu1Eb5v0SIVB-uDLYqQ,AAMC--8/narrativ

e-3.png

CHK@Akupw149t2lA50aMSDNIP0dvohsP47IhD9u-sMeDXXg,SLXJ55pld7Dfid17D4ViewxZ9ksUmeJtXmSbXg4rVUQ,AAMC--8/narrativ

e-4.png

CHK@NK30gKnbeqMPWI1QWoNL9nzdFVNKwXw0WoQShnDhgPM,KY7J9YNV3nlYygFfiFSZaIJFydfZp1UQIjkNrbgFdJw,AAMC--8/narrativ

e-5.png

reply1:

So, the advice is to disable peer count sharing?

reply2:

My guess based on this, I'm looking forward to reading the paper, is that they work on the simple basis of amount of requests going through the node.

I think their method makes sense to a certain degree. If you know the approximate distribution of peers and the number of peers a node has you can to some extent say how many blocks you would expect to see if that node requests a known file.

This would, in my reasoning, be offset by the possibility of several people requesting the file at the same time and sharing peers, but since the drop off of requests per hop is proportional to an exponential function of the number of peers each hop in the chain has it could probably paint a rather good picture.

A small thought experiment ignoring link length distribution and assuming even distribution of blocks (I assume we can ignore the link length distribution since it would only add some proportionality factor based on my own location and the peer's location.):

FMS Post 041917.txt

Node A: 50 peers
Node B: 40 peers
LEA: 20 peers

All three nodes are connected to each other.

If Node A requests a file with 100 blocks, ignoring link length distribution, both node B and LEA should expect to see about

$\text{Number File Blocks} / \text{Number peers} = 100 / 50 = 2$

requests from node A. LEA should then see about

$2 / 40 \sim 0$ (~ is approximately)

from node B that have been forwarded from Node A. LEA will probably see about 0 blocks forwarded from Node B since it has 40 peers to choose from when forwarding the blocks another step.

If we assume the following connection chain:

Node A -> Node B -> LEA, with the same peer count and ignoring link length distribution.

To make Node B appear to be requesting a file that Node A is requesting a would have to send:

$(\text{File block number} / \text{Node B peer count}) * \text{Node A peer count} * \text{Node B peer count} = 2 * 50 * 40 = 4\ 000$

requests, since 4 000 requests distributed evenly over Node A's peers would result in Node B getting about 80 requests to distribute evenly over its peers resulting in LEA getting the same result of about 2 blocks.

If we assume that Node B has 39 * Node A as neighbours and the LEA node (40 total neighbours), it would require pretty much all the

FMS Post 041917.txt

neighbours requesting the file at the same time to make it appear to the LEA node as if the file was being requested by Node B.

I'm only guessing, but if I were to attack Freenet I would reason something like this and then do some simulations with link length distribution and come up with a good metric. We'll see if I'm completely off when that paper lands.